




AL

INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(51) International Patent Classification ⁶ : G06F 12/00	A1	(11) International Publication Number: WO 00/08563 (43) International Publication Date: 17 February 2000 (17.02.00)
(21) International Application Number: PCT/US99/15830 (22) International Filing Date: 14 July 1999 (14.07.99) (30) Priority Data: 09/129,042 4 August 1998 (04.08.98) US (71) Applicant: THE BOARD OF GOVERNORS FOR HIGHER EDUCATION, STATE OF RHODE ISLAND AND PROVIDENCE PLANTATIONS [US/US]; 199 Promenade Street, Providence, RI 02908 (US). (72) Inventors: YANG, Qing; 81 West Wind Road, Wakefield, RI 02879 (US). HU, Yiming; Apartment G-5, Graduate Terrace, Kingston, RI 02881 (US). (74) Agents: STEVENS, Richard, L. et al.; Samuels, Gauthier & Stevens, Suite 3300, 225 Franklin Street, Boston, MA 02110 (US).		(81) Designated States: AL, AM, AT, AU, AZ, BA, BB, BG, BR, BY, CA, CH, CN, CU, CZ, DE, DK, EE, ES, FI, GB, GD, GE, GH, GM, HR, HU, ID, IL, IN, IS, JP, KE, KG, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, MD, MG, MK, MN, MW, MX, NO, NZ, PL, PT, RO, RU, SD, SE, SG, SI, SK, SL, TJ, TM, TR, TT, UA, UG, UZ, VN, YU, ZW, ARIPO patent (GH, GM, KE, LS, MW, SD, SL, SZ, UG, ZW), Eurasian patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European patent (AT, BE, CH, CY, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, GW, ML, MR, NE, SN, TD, TG). Published With international search report.

(54) Title: REDUNDANT, ASYMMETRICALLY PARALLEL DISK CACHE FOR A DATA STORAGE SYSTEM



(57) Abstract

A data storage system (90) includes a CPU (12), I/O bus (14), redundant write caches, a disk controller (16) and an array of disks (92-95). One of the redundant write caches is a primary write cache (20) of RAM or NVRAM, and another is a backup write cache (22) having a hybrid memory structure composed of a small amount of NVRAM with a cache-disk space mapped to disk. The cache-disk space may be located on a single disk or distributed over a plurality of disks within the disk array (92-95). In one embodiment, the array of disks is configured as a RAID architecture. The redundant write caches are asymmetric since the primary write cache (20) and the backup write cache (22) have different sizes and structures. The caching arrangement of the present invention has a significant reliability advantage over single NVRAM write caches, and provide a cost-effective architecture for very large write caches.

FOR THE PURPOSES OF INFORMATION ONLY

Codes used to identify States party to the PCT on the front pages of pamphlets publishing international applications under the PCT.

AL	Albania	ES	Spain	LS	Lesotho	SI	Slovenia
AM	Armenia	FI	Finland	LT	Lithuania	SK	Slovakia
AT	Austria	FR	France	LU	Luxembourg	SN	Senegal
AU	Australia	GA	Gabon	LV	Latvia	SZ	Swaziland
AZ	Azerbaijan	GB	United Kingdom	MC	Monaco	TD	Chad
BA	Bosnia and Herzegovina	GE	Georgia	MD	Republic of Moldova	TG	Togo
BB	Barbados	GH	Ghana	MG	Madagascar	TJ	Tajikistan
BE	Belgium	GN	Guinea	MK	The former Yugoslav Republic of Macedonia	TM	Turkmenistan
BF	Burkina Faso	GR	Greece			TR	Turkey
BG	Bulgaria	HU	Hungary	ML	Mali	TT	Trinidad and Tobago
BJ	Benin	IE	Ireland	MN	Mongolia	UA	Ukraine
BR	Brazil	IL	Israel	MR	Mauritania	UG	Uganda
BY	Belarus	IS	Iceland	MW	Malawi	US	United States of America
CA	Canada	IT	Italy	MX	Mexico	UZ	Uzbekistan
CF	Central African Republic	JP	Japan	NE	Niger	VN	Viet Nam
CG	Congo	KE	Kenya	NL	Netherlands	YU	Yugoslavia
CH	Switzerland	KG	Kyrgyzstan	NO	Norway	ZW	Zimbabwe
CI	Côte d'Ivoire	KP	Democratic People's Republic of Korea	NZ	New Zealand		
CM	Cameroon			PL	Poland		
CN	China	KR	Republic of Korea	PT	Portugal		
CU	Cuba	KZ	Kazakstan	RO	Romania		
CZ	Czech Republic	LC	Saint Lucia	RU	Russian Federation		
DE	Germany	LI	Liechtenstein	SD	Sudan		
DK	Denmark	LK	Sri Lanka	SE	Sweden		
EE	Estonia	LR	Liberia	SG	Singapore		

Title

Redundant, Asymmetrically Parallel Disk Cache For A Data Storage System

Technical Field

The invention relates to the field of redundant data storage systems, and in particular to a data storage system architecture having a redundant, asymmetrical disk cache.

Background of the Invention

A well known disk array system architecture is referred to as "redundant array of independent/inexpensive disks" (RAID). The RAID system architecture provides a large amount of data storage in a reasonably reliable manner. Several popular RAID system architectures are set forth in the paper entitled "*A Case for Redundant Arrays of Inexpensive Disks (RAID)*", Patterson et al., Proc. ACM SIGMOD, June 1988. These architectures include RAID-1, RAID-2, RAID-3, RAID-4 and RAID-5. U.S. Patent 5,526,482 entitled "*Storage Device Array Architecture With Copyback Cache*" briefly discusses each of these architectures, and in the interest of brevity a similar discussion will not be repeated herein.

In general, each of the RAID system architectures includes a plurality of disks that are controlled by a RAID disk controller. When a central processing unit (CPU) sends information to the disk controller for storage on disk, the controller directs how the information shall be stored on the plurality of disks to ensure that a failure of any one of the disks will not cause the system to loose information.

Modern disks are impressively reliable, with a Mean Time To Failure (MTTF) of up to about 1 million hours. Such a low failure rate, coupled with the inherent redundancy provided by the RAID system architecture provides a Mean Time To Data Loss (MTTDL) of around several hundreds of millions of hours in a typical RAID-5 system architecture. However, this is not the system MTTDL, since the reliability of the system is only as good as the reliability of its most unreliable component.

Many RAID systems employ a *single* cache in front of the plurality of disks. However, a single cache creates a single point failure mode which could disable the entire RAID system and cause the loss of data. Typical RAID caches use an NVRAM cache

(e.g., battery backed RAM) which has a MTTF of only about 15,000 hours. Therefore, failure of the NVRAM cache results in a significantly higher risk of data loss in comparison to the risk associated with a disk failure.

To overcome this reliability problem, some high-end RAID systems use *dual-copy* caches (i.e., a primary cache and a backup cache) so a failure in one cache does not cause the loss of data. That is, the other cache is intact which ensures the integrity of the data stored on disk. When a write request arrives from the CPU, the controller writes the data independently into the two caches. Although the use of dual-copy caches overcomes the reliability problem of the single-copy cache, the high cost of NVRAM makes dual-copy caches prohibitively expensive, particularly for large caches.

Therefore, there is a need for a redundant, inexpensive write cache for use in a data storage system, including RAID systems.

Summary of the Invention

An object of the present invention is to provide a data storage system which employs an inexpensive, redundant write cache.

Briefly, according to the present invention, a data storage system includes redundant write caches, a disk controller and an array of disks. One of the redundant write caches is a primary write cache and another is a backup write cache having a hybrid memory structure comprising a cache-disk space mapped to at least one of the disks.

The cache-disk space may be located on a single disk within the disk array, or distributed over a number of the plurality of disks in the array. In one embodiment, the array of disks can may be configured as a RAID architecture.

The data storage system of the present invention preferably employs a conventional, fast-write-fast-read primary write cache and a non-volatile, hybrid memory backup write cache. The backup write cache includes NVRAM, and small and random writes are buffered in a least recently used (LRU) cache located in the NVRAM. If the LRU cache is full, LRU data blocks in the LRU cache are transferred into a segment buffer which is also located in the backup cache NVRAM to make room for new write requests. When the segment buffer is full, its entire contents are written into the cache-disk space. Notably, combining the data from smaller writes into a block of data for a larger write reduces the number of smaller writes, and the segment buffer is quickly made available for

additional requests so the two level backup cache appears to a host CPU as a large NVRAM. As a result, the backup write cache may achieve the same write speed as the primary cache.

The redundant write caches are asymmetric since the primary write cache and the
5 backup write cache have different structures.

The system is relatively inexpensive since the amount NVRAM in the backup cache is relatively small, ranging from hundreds of KB to several MB, and the cost of the disk space for the cache-disk space is significantly less than a large amount of NVRAM.

Advantageously, the caching arrangement of the present invention has a significant
10 reliability advantage over conventional single NVRAM write caches, and a cost advantage over dual-copy NVRAM caches. In addition, the present invention provides a cost-effective architecture for very large write caches capable of masking the effects of small writes for high-end data storage systems that would otherwise have to use dual-copy, identical NVRAM caches.

15 These and other objects, features and advantages of the present invention will become apparent in light of the following detailed description of preferred embodiments thereof, as illustrated in the accompanying drawings.

Brief Description of the Drawings

Fig. 1 is a functional block diagram illustration of a data storage system
20 architecture according to the present invention;

Fig. 2 is a functional block diagram illustration of a backup write cache;

Fig. 3 is a flow chart illustration of the processing steps associated with the destaging process;

Fig. 4 is a functional block diagram illustration of an alternative embodiment data
25 storage system;

Fig. 5 is a pictorial illustration of the cache-disk memory space;

Fig. 6 illustrates a triple redundant data storage system; and

Fig. 7 illustrates yet another triple data storage system.

Detailed Description of the Invention

30 Fig. 1 illustrates a functional block diagram of a data storage system 10. The

system 10 communicates with a central processing unit (CPU) 12 via a bus 14 (e.g., a small computer systems interface (SCSI) bus). The system 10 includes a controller 16 which communicates with a read cache 18, a primary write cache 20 and a backup write cache 22. During normal operations (i.e., in the absence of a system failure), the primary
5 write cache 20 and the backup write cache 22 contain redundant data. As known, if either of these two caches fails, the data can be reconstructed from the other.

Each of the caches 18, 20, 22 reads data from or writes data to a plurality of disks 26-30. In a preferred embodiment, the plurality of disks 26-30 are configured as a redundant array of independent/inexpensive disks (RAID). However, one of ordinary skill
10 will recognize that the redundant cache architecture of the present invention is applicable to virtually any input/output (I/O) system.

The primary write cache 20 generally employs NVRAM for storage (e.g., battery backed RAM) to ensure data is retained in the event of a power failure (or other system failures). However, to reduce system cost the primary write cache 20 may employ DRAM
15 rather than NVRAM. In contrast, according to the present invention, the backup write cache 22 has a hierarchical, hybrid-memory structure.

Fig. 2 illustrates a functional block diagram of the backup write cache 22. The cache has a hybrid memory structure that includes NVRAM 40, volatile RAM 41 (preferably DRAM) and a cache-disk space 42. The size of the NVRAM 40 ranges from
20 several hundred kB to several MB in size, while the cache-disk space 42 has the capacity to store about several hundred MB. The sizes discussed herein are used only for their relative sizes, and one of ordinary skill will recognize that the present invention is not limited to the sizes discussed herein.

The backup cache NVRAM 40 is partitioned to provide a last recently used (LRU)
25 cache 43, a plurality of segment buffers 44, 46 and a hash table 48. The recently accessed data resides in the LRU cache 43. The less recently accessed data is kept in the cache-disk space 42, and is preferably organized into a plurality of disk segments 56-60 similar to a log-structured file system such as Sprite LFS and the BSD LFS. Each disk segment 56-60 contains a number of slots each of which can hold one data block. In one embodiment,
30 each block is 32 kb. Data blocks stored in the segments 56-60 are addressed by their segment IDs and slot IDs. Data blocks stored in the LRU cache 43 are addressed by their

logical block addresses (LBAs).

The hash table 48 contains location information for each of the valid data blocks in the backup cache 22. Specifically, the location information identifies whether a block is in the LRU cache 43 or in the cache-disk space 42, as well as the address of the
5 information in the LRU cache 43 or the cache-disk space 42.

The backup cache 22 also includes a disk segment table 61 which is preferably located in the volatile RAM 41. The details of the disk segment table 61 shall be discussed herein in detail.

Referring to Figs. 1 and 2, the cache-disk space 42 is preferably physically located
10 on a portion of the disk 30. Significantly, assigning the cache-disk space 32 to disk advantageously implements the cache-disk space 42 without additional system hardware cost, since the plurality of disks 26-30 may include one or more spare disks that are available in the event an active disk fails. In many prior art systems, the spare disks are not used during normal operation and therefore, it is difficult to detect if one of the spares
15 has failed. However, by mapping the cache-disk space 42 to a spare disk (e.g., disk 30) it is easier to detect when the spare disk is working properly, thus reducing the possibility of a latent failure of a spare disk.

When the storage system 10 receives a write request from the CPU 12, the disk controller 16 invalidates any data copy in the read cache 18, and sends the data to be
20 written to both the primary write cache 20 and the LRU cache 43 of the backup write cache 22. If there is space available in the write caches 20, 22, the data is copied to the available space immediately. To indicate that the data block is located in the backup LRU cache 43, the controller 16 creates a hash entry which is stored in the hash table 48. Once the data is written into both the primary write cache 20 and the backup write cache 22, the
25 controller 16 sends an acknowledgment to the CPU 12 signaling that the write request is complete.

If there is no space left in the primary write cache 20, the controller 16 first tries to discard a "clean" block of data from the cache 20 to make room for the new request.

A "clean" block of data is defined as a cache resident block of data that has already been
30 successfully copied to disk. However, if the controller 16 cannot find a "clean" block to discard, it executes a foreground destaging process (to be described in detail hereinafter)

which writes a least-recently-used data block to the array of disks 24. When the least-recently-used block is safely written into the array of disks 24, the space that was used for the least-recently-used data block in the primary write cache 20 is freed for the incoming request. The controller 16 also invalidates the associated data in the backup write cache 5 22.

If the backup cache 22 is full (i.e., the LRU cache 43 is full), the controller 16 selects an empty one of the segment buffers 44, 46 and designates it as the "current" segment buffer. An LRU cache resident data block is then copied to the "current" segment buffer and the corresponding entries in the hash table 48 and the disk segment table 50 are 10 modified to indicate that the data block is now in the "current" segment buffer. The disk segment table 50 contains information regarding which slots and disk segments are free, and can be used to cache data. Space within the LRU cache 43 is then available to accept incoming requests. Subsequent write requests may continue to evict least-recently-used blocks from the LRU cache 43 and transfer the blocks to the "current" segment buffer until 15 it is full. Once the "current" segment buffer is full, the controller 16 designates another segment buffer as the "current" segment buffer and continues to accept data from the LRU cache 43. The controller then writes the contents of the full segment buffer into the cache-disk space 42, preferably in one large write during background.

Since the full segment buffer is written to the cache-disk space 42 in one large 20 write instead of many small writes, the full segment buffer is quickly again made available when the disk write finishes. Therefore, the NVRAM 40 and the cache-disk space 42 of the backup write cache 22 collectively appear to the controller 16 as a large uniform NVRAM write cache, whose write performance is analogous to the primary write cache 20.

25 The size of the NVRAM resident segment buffers 44, 46 directly affects the write efficiency. That is, the larger the segment buffers 44, 46, the smaller the write overheads caused by disk seeking and rotational latencies. However, for a given NVRAM size, as the size of the segment buffers 44, 46 increases, the smaller the LRU cache 43 becomes since the segment buffers 44, 46 and the LRU cache 43 are both physically resident in 30 NVRAM. Therefore, there is a trade-off between segment buffer size and LRU cache size.

When the controller 16 receives a read request from the CPU 12, the read cache

18 and the primary write cache 20 are searched. If there is a cache hit, the data is immediately read. However, if there is a cache miss, the LRU block (not shown) in the read cache 18 is discarded and its buffer space is freed. The requested data is then read from the array of disks 24 and stored in the freed LRU block of the read cache 18. The backup cache 22 is generally not involved in read operations.

As discussed briefly discussed above, the destaging process involves moving "dirty" data resident in the primary write cache 20 to disk. Dirty data is defined as data that is stored in the primary write cache, but not yet stored in disk. Destaging is performed in background when the controller 16 detects an idle period, or when the number of "dirty" blocks in the primary write cache 20 exceeds a predetermined value (i.e., a high water mark).

Fig. 3 is a flow chart illustration of processing steps 70 associated with the destaging process. This process 70 is executed by the controller regularly in background, and on an as-needed-basis in foreground. Upon entry to the process 70, the controller 16 (Fig. 1) executes step 72 to identify "dirty" blocks of data resident in the primary write cache 20 (Fig. 1). Step 74 is then executed to read the old data and the parity value associated with the old data. The old data and the old parity value are read from disk. A new parity value based upon the identified "dirty" data and the old data is then computed in step 76, and the new data and parity value are written to the plurality of disks 24 in step 78. To indicate that the data is now resident in disk, step 80 is performed to designate the blocks of data that were written to disk in step 78 as "clean". The controller 16 (Fig. 1) then executes step 82 to invalidate the copy of the data stored in the backup write cache 22 (Fig. 1) that is identical to the data written to disk in step 78.

If the data to be invalidated in the backup write cache 22 is resident in the NVRAM LRU buffer 43, the controller releases the LRU buffer 43 as necessary. However, if the data to be invalidated is resident in one of the segment buffers 44, 46, or one of the disk segments 56-60, the controller 16 invalidates the data by marking the corresponding segment slot stored in the disk segment table 61 (Fig. 2) as "invalid". The controller 16 also deletes the corresponding hash entry resident in the hash table 48.

During background the destaging process 70 preferably runs continuously until the controller processing returns to foreground. Alternatively, the destaging process may only

run continuously during background until the number of dirty blocks in the primary cache falls below the predetermined low water-mark. During background the controller may also execute known tasks such as built-in-test (BIT).

During foreground the destaging process 70 is called when the controller 16 (Fig. 1) receives a new write request and the entire primary write cache 20 is occupied by "dirty" data. This may occur when a large write overflows the primary write cache 20. In this case, the destaging process 100 (Fig. 3) is initiated to free space for the incoming request.

Notably, the backup cache is not read or written to during either the background or foreground destaging process. Therefore the slower read speed the cache-disk will not affect the destaging performance.

Fig. 4 illustrates a functional block diagram of an alternative embodiment data storage system 90. This alternative embodiment system 90 is substantially the same as the system 10 illustrated in Fig. 1, with the principal exception that the cache-disk space 42 (Fig. 2) is mapped into a plurality of disks 92-95, rather than one disk (e.g., disk 30 - Fig. 1). Specifically, the cache-disk space 42 is logically partitioned into sub-spaces that are distributed over the plurality of disks 92-95.

Referring to Fig. 5, the cache-disk space 42 includes a plurality of sub-spaces 42a-42d that are mapped into the plurality of the disks. For example, the first sub-space 42a which includes a first plurality of disk segments 56-58, is mapped to disk #1 92 (Fig. 4).

The second sub-space 42b includes a second plurality of disk segments 98-100 and is mapped to disk #2 93 (Fig. 4). Similarly, cache disk sub-space 42d includes a plurality of segments 102-104 and is mapped to disk #N 95 (Fig. 4). As a result, the cache-disk space 42 is subdivided and mapped into the plurality of disks 93-95 (Fig. 4).

The controller 16 may automatically reconfigure the system 10 illustrated Fig. 1 to the system 90 in Fig. 4 if the spare disk 30 (Fig. 1) becomes an active disk due to a disk failure. Specifically, the controller reconfigures the backup write cache 42 to operate using a logical cache-disk space that is partitioned into subspaces and mapped into the active disks since the dedicated spare disk may no longer be available. When using the *logical* cache-disk space, the data written into the cache partitions on the disks do not involve in parity operations.

For a system with a logical cache-disk (rather than a dedicated physical cache disk), the segment buffer size can not be too large because segment writes must compete with normal reads in data disks. Large segment buffer sizes may result in lower read performance. Computer based simulations have shown that using four 128 KB segment
5 buffers achieves good system performance.

Fig. 6 illustrates another alternative embodiment data storage system 110. This data storage system is triple redundant system. Specifically, the system includes a primary cache 112 and two backup caches 114,116. During normal operation the primary cache 112 and the two backup caches 114,116 provide triple redundancy protection. Notably,
10 since the back-up caches 114,116 include non-volatile memory, they provide dual redundancy in the event of a power failure. In this embodiment, dedicated cache disks 118, 120 cooperate with the backup write caches 114,116 respectively.

Fig. 7 illustrates yet another alternative embodiment data storage system 130. This system 130 is substantially the same as the system 110 illustrated in Fig. 7, with the
15 principal exception that the cache-disk space is mapped into a plurality of disks 132-135.

Although one destaging process has been described, one of ordinary skill will recognize that other techniques such as a known linear threshold scheduling algorithm may also be used.

Although the present invention has been shown and described with respect to
20 several preferred embodiments thereof, various changes, omissions and additions to the form and detail thereof, may be made therein, without departing from the spirit and scope of the invention.

What is claimed is:

- 1 1. A data storage system which receives data from and provides data to a central
2 processing unit (CPU), the system comprising:
 - 3 a disk controller which receives data from the CPU;
 - 4 a plurality of disks;
 - 5 a primary write cache which receives write data to be written to said plurality of
6 disks; and
 - 7 a backup write cache which also receives said write data and writes said data to
8 backup cache disk space located on at least one of said plurality of disks.
- 1 2. The system of claim 1, wherein said backup write cache comprises:
 - 2 NVRAM which is sub-divided to provide an LRU buffer, a hash table and a
3 plurality of segment buffers.
- 1 3. The system of claim 2, wherein said primary write cache includes significantly
2 more RAM than said backup write cache.
- 1 4. The system of claim 1, wherein said backup cache disk space is located on one of
2 said plurality of disks.
- 1 5. The system of claim 1, wherein said cache-disk space is located on more than one
2 of said plurality of disks.
- 1 6. The system of claim 1, wherein said backup write cache further comprises a disk
2 segment table.
- 1 7. The system of claim 6, wherein said disk segment table is located in a random array
2 of memory separate from said LRU cache.
- 1 8. The system of claim 7, wherein said disk segment table is stored in DRAM.
- 1 9. The system of claim 3, wherein said primary write cache RAM comprises
2 NVRAM.
- 1 10. The system of claim 3, wherein said primary write cache RAM comprises DRAM.
- 1 11. A redundant, asymmetrical disk cache for a data storage system that includes a disk

2 controller and at least one disk, said disk cache comprising:

3 a primary write cache which receives data to be written to the disks; and

4 a backup write cache which also receives the data to be written to disk, wherein
5 said backup write cache includes a hybrid memory structure comprising a cache-disk space
6 which is mapped into the disk.

1 12. The data storage system of claim 11, wherein said backup write cache further
2 comprises NVRAM that is partitioned to provide an LRU cache, a hash value table and a
3 plurality of segment tables.

1 13. The data storage system of claim 11, wherein said cache-disk space is partitioned
2 into subspaces and mapped into a plurality of disks.

1 14. A backup write cache for a data storage system having at least one disk, the backup
2 write cache comprising:

3 non-volatile random access memory which receives information to be written to the
4 disk; and

5 a cache-disk space which is mapped into the disk.

1 / 7




FIG. 1




FIG. 2

3 / 7

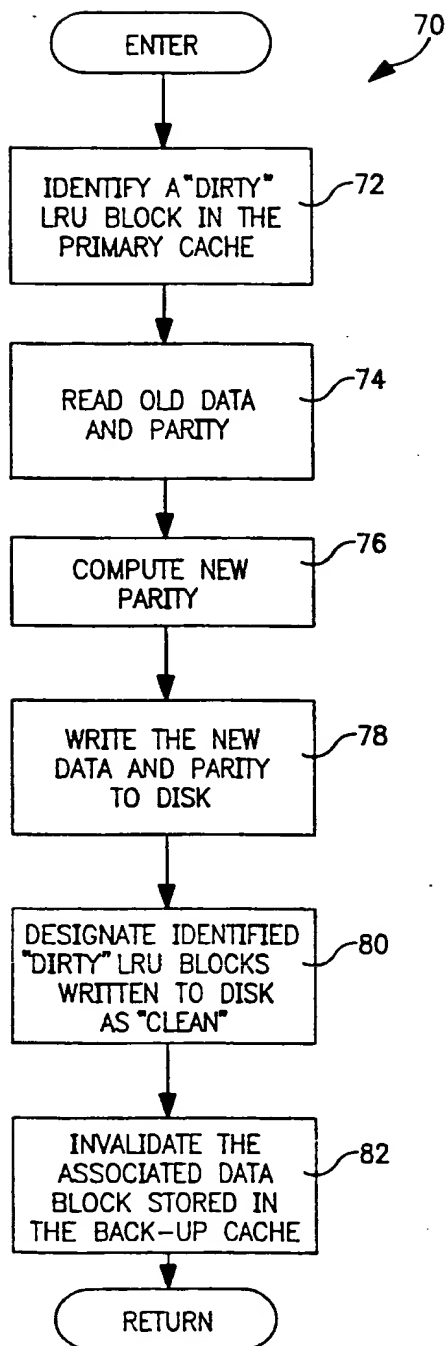


FIG. 3




FIG. 4




FIG. 5




FIG. 6

7 / 7




FIG. 7

INTERNATIONAL SEARCH REPORT

International application No.

PCT/US99/15830

A. CLASSIFICATION OF SUBJECT MATTER

IPC(6) :G06F 12/00

US CL :711/119,162

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)

U.S. : 711/105,118,119,159,160,162,

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practicable, search terms used)

APS: multiple cache, backup, primary, secondary, redundant, LRU, hash, table, NVRAM, non-volatile, plurality

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X --- Y	US 5,526,482 A (STALLMO ET AL.) 11 JUNE 1996, col. 6, lines 17-45; col. 18, lines 1-11; col. 16, lines 50-61; Fig. 4; col. 11, lines 1-4; col. 12, lines 23-25;	1, 4, 14 ----- 2-3, 5-10
X --- Y	US 5,588,129 A (BALLARD) 24 DECEMBER 1996, col 3, lines 7-61; col 6, lines 29-31; col. 9, lines 21-25; col. 10, lines 48-49; col. 11, lines 34-36; col. 7, lines 22-25;	11-13 ----- 5-8
Y	US 5,577,226 A (PERCIVAL) 19 NOVEMBER 1996, col. 22, lines 58-60; col. 14, 41-43;	2-3, 9-10
Y	US 5,535,381 A (KOPPER) 09 JULY 1996, col. 3, lines 48-53;	3, 9-10

☒ Further documents are listed in the continuation of Box C. ☐ See patent family annex.

* Special categories of cited documents:	*T	later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention
A document defining the general state of the art which is not considered to be of particular relevance	*X*	document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone
E earlier document published on or after the international filing date	*Y*	document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art
L document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)	*A*	document member of the same patent family
O document referring to an oral disclosure, use, exhibition or other means		
P document published prior to the international filing date but later than the priority date claimed		

Date of the actual completion of the international search

26 AUGUST 1999

Date of mailing of the international search report

19 OCT 1999

Name and mailing address of the ISA/US
Commissioner of Patents and Trademarks
Box PCT
Washington, D.C. 20231

Facsimile No. (703) 305-3230

Authorized officer

Brian R. Peugh

Telephone No. (703) 306-5843

Joni Hill

INTERNATIONAL SEARCH REPORT

International application No.

PCT/US99/15830

C (Continuation). DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
Y	US 5,341,493 A (YANAI ET AL.) 23 AUGUST 1994, col. 1, lines 42-45;	8, 10
A	US 5,787,466 A (BERLINER) 28 JULY 1998, abstract; col. 7, lines 14-60;	1-14